



USPTO

[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: ☒ The ACM Digital Library ☐ The Guide



THE ACM DIGITAL LIBRARY


[Feedback](#) [Report a problem](#) [Satisfaction survey](#)
Terms used service groups fail over

Found 4 of 447 searched out of 447.

Sort results by

[Save results to a Binder](#)

Display results

[Search Tips](#)
☐ Open results in a new window
Try an [Advanced Search](#)Try this search in [The ACM Guide](#)

Results 1 - 4 of 4

Relevance scale ☐ ☐ ☐ ☐ ☐

1 [A load cluster management system using SNMP and web](#)

Myung-Sup Kim, Mi-Joung Choi, James W. Hong

November 2002 **International Journal of Network Management**, Volume 12 Issue 6**Publisher:** John Wiley & Sons, Inc.Full text available: [pdf\(355.47 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Clustered servers for Internet service is a popular solution to cope with the explosive increase in client requests. The high probability of service failure in cluster servers make the cluster management system necessary to provide high availability and convenient administrator control. In this paper, we present the design and implementation of a load cluster management system (LCMS) based on SNMP and Web technology. Our LCMS implementation has been deployed on a commercial ultra-dense server.

2 [Lessons learned from the novell and groupWise upgrade of the summer 2003 or "what we did on our summer vacation!"](#)

Robert L. Barley, Yancy Phillips

October 2004 **Proceedings of the 32nd annual ACM SIGUCCS conference on User services****Publisher:** ACM PressFull text available: [pdf\(174.72 KB\)](#) Additional Information: [full citation](#), [abstract](#), [index terms](#)

This presentation will focus on the effects a recent major hardware and software upgrade had on the Indiana State University end user community and how collaboration between Technical and User Services resulted in a positive experience for the end user community. There were many hours of overtime put into this project and lots of concerns from User Services personnel about having to go out and "touch" each computer to facilitate necessary changes. The effects of the unsolicited visits produce ...

Keywords: novell upgrades, reliability, standardization, support

3 [Operating and runtime systems for high-end computing systems: MOLAR: adaptive runtime support for high-end computing operating and runtime systems](#)

Christian Engelmann, Stephen L. Scott, David E. Bernholdt, Narasimha R. Gottumukkala, Chokchai Leangsuksun, Jyothish Varma, Chao Wang, Frank Mueller, Aniruddha G. Shet, P. Sadayappan

April 2006 **ACM SIGOPS Operating Systems Review**, Volume 40 Issue 2**Publisher:** ACM Press

Full text available:  [pdf\(522.07 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

MOLAR is a multi-institutional research effort that concentrates on adaptive, reliable, and efficient operating and runtime system (OS/R) solutions for ultra-scale high-end scientific computing on the next generation of supercomputers. This research addresses the challenges outlined in FAST-OS (forum to address scalable technology for runtime and operating systems) and HECRTF (high-end computing revitalization task force) activities by exploring the use of advanced monitoring and adaptation to i ...

Keywords: RAS, availability, fault tolerance, group membership, high-end computing, monitoring, reliability

4 [Can unstructured P2P protocols survive flash crowds?](#)



Dan Rubenstein, Sambit Sahu

June 2005 **IEEE/ACM Transactions on Networking (TON)**, Volume 13 Issue 3

Publisher: ACM Press

Full text available:  [pdf\(345.18 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Today's Internet periodically suffers from hot spots, a.k.a., flash crowds. A hot spot is typically triggered by an unanticipated news event that triggers an unanticipated surge of users that request data objects from a particular site, temporarily overwhelming the site's delivery capabilities. During this time, the large majority of users that attempt to get these objects face the frustrating experience of not being able to retrieve the content they want while still being able to communicate ef ...

Keywords: average case analysis, flash crowds, peer-to-peer

Results 1 - 4 of 4

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)



USPTO

[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: ☒ The ACM Digital Library ☐ The Guide

"fail over" +lock

SEARCH

THE ACM DIGITAL LIBRARY


[Feedback](#) [Report a problem](#) [Satisfaction survey](#)
Terms used **fail over lock**Found **33** of **160** searched out of **160**.

Sort results by

relevance

Display results

expanded form

[Save results to a Binder](#) [Search Tips](#)☐ Open results in a new window[Try an Advanced Search](#)[Try this search in The ACM Guide](#)

Results 1 - 20 of 33

Result page: **1** [2](#) [next](#)Relevance scale ☐ ☐ ☐ ☐ ☐

1 [Manageability, availability, and performance in porcupine: a highly scalable, cluster-based mail service](#)



Yasushi Saito, Brian N. Bershad, Henry M. Levy

August 2000 **ACM Transactions on Computer Systems (TOCS)**, Volume 18 Issue 3

Publisher: ACM Press

Full text available: pdf(2.52 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

This paper describes the motivation, design and performance of Porcupine, a scalable mail server. The goal of Porcupine is to provide a highly available and scalable electronic mail service using a large cluster of commodity PCs. We designed Porcupine to be easy to manage by emphasizing dynamic load balancing, automatic configuration, and graceful degradation in the presence of failures. Key to the system's manageability, availability, and performance is that sessions, data, and underlying ...

Keywords: cluster, distributed systems, email, group membership protocol, load balancing, replication

2 [Some Deadlock Properties of Computer Systems](#)



Richard C. Holt

September 1972 **ACM Computing Surveys (CSUR)**, Volume 4 Issue 3

Publisher: ACM Press

Full text available: pdf(1.46 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

3 [Research papers: correctness and trust: Middleware based data replication providing snapshot isolation](#)





Yi Lin, Bettina Kemme, Marta Patiño-Martínez, Ricardo Jiménez-Peris

June 2005 **Proceedings of the 2005 ACM SIGMOD international conference on Management of data**

Publisher: ACM Press

Full text available: pdf(470.92 KB) Additional Information: [full citation](#), [abstract](#), [references](#)

Many cluster based replication solutions have been proposed providing scalability and fault-tolerance. Many of these solutions perform replica control in a middleware on top of the database replicas. In such a setting concurrency control is a challenge and is often performed on a table basis. Additionally, some systems put severe requirements on transaction programs (e.g., to declare all objects to be accessed in advance). This paper addresses these issues and presents a middleware-based replica ...


-  Frontmatter (TOC, Letters, Philosophy of computer science, Interviewers needed, Taking software requirements creation from folklore to analysis, SW components and product lines: from business to systems and technology, Software engineering survey) 

September 2005 **ACM SIGSOFT Software Engineering Notes**, Volume 30 Issue 5

Publisher: ACM Press

Full text available:  pdf(1.98 MB) Additional Information: [full citation](#), [index terms](#)

- 5 Frangipani: a scalable distributed file system 

 Chandramohan A. Thekkath, Timothy Mann, Edward K. Lee

October 1997 **ACM SIGOPS Operating Systems Review , Proceedings of the sixteenth ACM symposium on Operating systems principles SOSP '97**, Volume 31 Issue 5

Publisher: ACM Press

Full text available:  pdf(2.20 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

- 6 High Availability Cluster Checklist 

Tim Burke

November 2000 **Linux Journal**

Publisher: Specialized Systems Consultants, Inc.

Full text available:  html(17.23 KB) Additional Information: [full citation](#), [abstract](#), [index terms](#)

With a variety of clustering services on the market, the ability to determine how well options meet your specific business needs is necessary.

- 7 Q focus: databases: Beyond relational databases 

 Margo Seltzer

April 2005 **Queue**, Volume 3 Issue 3

Publisher: ACM Press

Full text available:  pdf(654.63 KB)  html(35.21 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

There is more to data access than SQL.


- 8 Queue Focus: Game Development: Massively Multiplayer Middleware 

 Michi Henning

February 2004 **Queue**, Volume 1 Issue 10

Publisher: ACM Press

Full text available:  pdf(2.34 MB)  html(28.00 KB) Additional Information: [full citation](#), [index terms](#)

- 9 Potpourri: Fast and transparent recovery for continuous availability of cluster-based servers 

 Rosalia Christodouloupoulou, Kaloian Manassiev, Angelos Bilas, Cristiana Amza

March 2006 **Proceedings of the eleventh ACM SIGPLAN symposium on Principles and practice of parallel programming PPOPP '06**

Publisher: ACM Press

Full text available:  pdf(111.02 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Recently there has been renewed interest in building reliable servers that support continuous application operation. Besides maintaining system state consistent after a failure, one of the main challenges in achieving continuous operation is to provide fast reconfiguration. The complexity of the failure reconfiguration mechanisms employed and their overheads depend on the type of platform that is being used as a server and the

types of applications that need to be supported. In this paper we foc ...

Keywords: availability, distributed shared memory, fast failure reconfiguration, fault tolerance, scalability

10 A new approach to developing and implementing eager database replication protocols



Bettina Kemme, Gustavo Alonso

September 2000 **ACM Transactions on Database Systems (TODS)**, Volume 25 Issue 3

Publisher: ACM Press

Full text available: pdf(449.43 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Database replication is traditionally seen as a way to increase the availability and performance of distributed databases. Although a large number of protocols providing data consistency and fault-tolerance have been proposed, few of these ideas have ever been used in commercial products due to their complexity and performance implications. Instead, current products allow inconsistencies and often resort to centralized approaches which eliminates some of the advantages of replication. As an ...

Keywords: database replication, fault-tolerance, group communication, isolation levels, one-copy-serializability, replica control, total error multicast

11 Affinity-based management of main memory database clusters



Minwen Ji

November 2002 **ACM Transactions on Internet Technology (TOIT)**, Volume 2 Issue 4

Publisher: ACM Press

Full text available: pdf(553.96 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We study management strategies for main memory database clusters that are interposed between Internet applications and back-end databases as content caches. The task of management is to allocate data across individual cache databases and to route queries to the appropriate databases for execution. The goal is to maximize effective cache capacity and to minimize synchronization cost. We propose an affinity-based management system for main memory database cLusters (*ALBUM*). *ALBUM* executes ea ...

Keywords: Main memory database, clustering, database administration, database cluster, file organization, query affinity, scalability

12 AFS—a secure distributed filesystem, Part III



Alf Wachsman

April 2005 **Linux Journal**, Volume 2005 Issue 132

Publisher: Specialized Systems Consultants, Inc.

Full text available: html(22.05 KB)

Additional Information: [full citation](#), [abstract](#), [index terms](#)

Reconfigure servers without changing mount points on the clients with this Kerberos-authenticated network filesystem.

13 Versatility and Unix semantics in namespace unification



Charles P. Wright, Jay Dave, Puja Gupta, Harikesavan Krishnan, David P. Quigley, Erez Zadok, Mohammad Nayyer Zubair

February 2006 **ACM Transactions on Storage (TOS)**, Volume 2 Issue 1

Publisher: ACM Press

Full text available: pdf(317.82 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Administrators often prefer to keep related sets of files in different locations or media, as it is easier to maintain them separately. Users, however, prefer to see all files in one

location for convenience. One solution that accommodates both needs is virtual namespace unification---providing a merged view of several directories without physically merging them. For example, namespace unification can merge the contents of several CD-ROM images without unpacking them, merge binary directories fr ...

Keywords: Namespace management, directory merging, snapshotting, stackable file systems, unification

14 Research sessions: consistency and availability: Highly available, fault-tolerant, parallel dataflows



Mehul A. Shah, Joseph M. Hellerstein, Eric Brewer

June 2004 **Proceedings of the 2004 ACM SIGMOD international conference on Management of data**

Publisher: ACM Press

Full text available: pdf(210.17 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citing](#)

We present a technique that masks failures in a cluster to provide high availability and fault-tolerance for long-running, parallelized dataflows. We can use these dataflows to implement a variety of continuous query (CQ) applications that require high-throughput, 24x7 operation. Examples include network monitoring, phone call processing, click-stream processing, and online financial analysis. Our main contribution is a scheme that carefully integrates traditional query processing techniques for ...

15 Poster session: papers included: System centric mission critical simulation model for MES automation

Amit Jindal, Rajkumar Khandelwal

December 2005 **Proceedings of the 37th conference on Winter simulation WSC '05**

Publisher: Winter Simulation Conference

Full text available: pdf(217.40 KB) Additional Information: [full citation](#), [abstract](#), [references](#)

This paper describes the system centric simulation methodology used for stress testing of Manufacturing Execution System (MES) in Intel. System centric simulation involves testing such that the system components (infrastructure stack and software) are characterized for the load they would experience in production, irrespective of how that load is exerted. A new manufacturing execution system software is introduced in Intel's latest fabrication facility. Validation of the product under stress is ...

16 Manageability, availability and performance in Porcupine: a highly scalable, cluster-based mail service



Yasushi Saito, Brian N. Bershad, Henry M. Levy

December 1999 **ACM SIGOPS Operating Systems Review , Proceedings of the seventeenth ACM symposium on Operating systems principles SOSP '99**, Volume 33 Issue 5

Publisher: ACM Press

Full text available: pdf(1.62 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citing](#), [index terms](#)

This paper describes the motivation, design, and performance of Porcupine, a scalable mail server. The goal of Porcupine is to provide a highly available and scalable electronic mail service using a large cluster of commodity PCs. We designed Porcupine to be easy to manage by emphasizing dynamic load balancing, automatic configuration, and graceful degradation in the presence of failures. Key to the system's manageability, availability, and performance is that sessions, data, and underlying serv ...

17 Database replication with Slony-I

Ludovic Marcotte

June 2005 **Linux Journal**, Volume 2005 Issue 134

Publisher: Specialized Systems Consultants, Inc.

Full text available: html(22.00 KB) Additional Information: [full citation](#), [abstract](#), [index terms](#)

Move up to a highly available cluster without leaving behind the open-source database you trust.

18 Session 7: OS architecture II: Increasing relevance of memory hardware errors: a case for recoverable programming models



Dejan Milojevic, Alan Messer, James Shau, Guangrui Fu, Alberto Munoz

September 2000

Proceedings of the 9th workshop on ACM SIGOPS European workshop: beyond the PC: new challenges for the operating system

Publisher: ACM Press

Full text available: pdf(99.65 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

It is a common belief that most of computer system failures nowadays stem from programming errors. Computer systems are becoming more complex and harder to maintain and administer, making software errors an even more common case, while contemporary computer architectures are optimized for price and performance and not for availability. In this paper, we raise a case for an increasing relevance of memory hardware soft-errors. In particular with the introduction of 64-bit processors, memory scalin ...

19 Experience reports: software architecture II: The co-evolution of a hype and a software architecture: experience of component-producing large-scale EJB early adopters

Lutz Prechelt

May 2003

Proceedings of the 25th International Conference on Software Engineering

Publisher: IEEE Computer Society

Full text available: pdf(367.70 KB)



[Publisher Site](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

abaXX.components was one of the first API software products fully based on Enterprise JavaBeans™ (EJB) technology. We describe the evolution of its architecture as it moved from simply taking the initial EJB hype for the truth, through several intermediate stages, to using EJB simply as one of several encapsulated implementation techniques. So far, the public perception of how to use EJB properly evolved along a similar path, lagging 6 to 12 months behind.

20 Posters: Data versioning techniques for internet transaction management



Ramkrishna Chatterjee, Gopalan Arun

May 2005

Special interest tracks and posters of the 14th international conference on World Wide Web

Publisher: ACM Press

Full text available: pdf(134.78 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

An Internet transaction is a transaction that involves communication over the Internet using standard Internet protocols such as HTTPS. Such transactions are widely used in Internet-based applications such as e-commerce. With the growth of the Internet, the volume and complexity of Internet transactions are rapidly increasing. We present data versioning techniques that can reduce the complexity of managing Internet transactions and improve their scalability and reliability. These techniques have ...

Keywords: internet transaction, scalability, versioning

Results 1 - 20 of 33

Result page: [1](#) [2](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: ☒ The ACM Digital Library ☐ The Guide

THE ACM DIGITAL LIBRARY


[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

 Terms used **fail over lock**

Found 33 of 160

Sort results by


[Save results to a Binder](#)
[Try an Advanced Search](#)

Display results


[Search Tips](#)
[Try this search in The ACM Guide](#)
☐ Open results in a new window

Results 21 - 33 of 33

 Result page: [previous](#) [1](#) [2](#)

 Relevance scale ☐ ☐ ☐ ☐ ☐

21 [Systems and prototypes: Phoenix project: fault-tolerant applications](#)



Roger Barga, David Lomet

 June 2002 **ACM SIGMOD Record**, Volume 31 Issue 2

Publisher: ACM Press

 Full text available: pdf(847.56 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

After a system crash, databases recover to the last committed transaction, but applications usually either crash or cannot continue. The Phoenix purpose is to enable application state to persist across system crashes, transparent to the application program. This simplifies application programming, reduces operational costs, masks failures from users, and increases application availability, which is critical in many scenarios, e.g., e-commerce. Within the Phoenix project, we have explored how to ...

22 [Industrial session: potpourri: Recovery principles of MySQL Cluster 5.1](#)

Mikael Ronström, Jonas Orelund

 August 2005 **Proceedings of the 31st international conference on Very large data bases VLDB '05**

Publisher: VLDB Endowment

 Full text available: pdf(299.62 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

MySQL Cluster is a parallel main memory database. It is using the normal MySQL software with a new storage engine NDB Cluster. MySQL Cluster 5.1 has been adapted to also handle fields on disk. In this work a number of recovery principles of MySQL Cluster had to be adapted to handle very large data sizes. The article presents an efficient algorithm for synchronizing a starting node with very large data sets. It provides reasons for the unorthodox choice of a no-steal algorithm in the buffer manag ...

23 [A novel Fuzzy Logic Controller \(FLC\) for shortening the TCP channel roundtrip time by eliminating user buffer overflow adaptively](#)

Wilfred W. K. Lin, Allan K. Y. Wong, Tharam S. Dillon

 January 2005 **Proceedings of the Twenty-eighth Australasian conference on Computer Science - Volume 38 ACSC '05**

Publisher: Australian Computer Society, Inc.

 Full text available: pdf(951.23 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

The proposed Fuzzy Logic Controller (FLC) is a novel approach for dynamic buffer tuning at the user/server level. It eliminates buffer overflow by ensuring that the buffer length always cover the queue length adaptively. The FLC and the AQM (active queue management) mechanisms at the router/system level together form a unified solution to stifle TCP (*Transmission Control Protocol*) channel buffer overflow over the Internet. The FLC contributes to: a) prevent the AQM resources dished out at ...

Keywords: FLC, PIDC, active queue management, buffer overflow, dynamic buffer tuning, internet, traffic pattern

24 Cryptography as an operating system service: A case study



Angelos D. Keromytis, Jason L. Wright, Theo De Raadt, Matthew Burnside

February 2006 **ACM Transactions on Computer Systems (TOCS)**, Volume 24 Issue 1

Publisher: ACM Press

Full text available: pdf(669.12 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Cryptographic transformations are a fundamental building block in many security applications and protocols. To improve performance, several vendors market hardware accelerator cards. However, until now no operating system provided a mechanism that allowed both uniform and efficient use of this new type of resource. We present the OpenBSD Cryptographic Framework (OCF), a service virtualization layer implemented inside the operating system kernel, that provides uniform access to accelerator functio ...

Keywords: Encryption, authentication, cryptographic protocols, digital signatures, hash functions

25 An open architecture for next-generation telecommunication services



Gregory W. Bond, Eric Cheung, K. Hal Purdy, Pamela Zave, J. Christopher Ramming

February 2004 **ACM Transactions on Internet Technology (TOIT)**, Volume 4 Issue 1

Publisher: ACM Press

Full text available: pdf(237.24 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

An open (in the sense of extensible and programmable) architecture for IP telecommunications must be based on a comprehensive strategy for managing feature interaction. We describe our experience with BoxOS, an IP telecommunication platform that implements the DFC technology for feature composition. We present solutions to problems, common to all efforts in IP telecommunications, of feature distribution, interoperability, and media management. We also explain how BoxOS addresses many deficiencies ...

Keywords: Component architectures, Intelligent Network architecture, Session Initiation Protocol, electronic mail, feature interaction, instant messaging, multimedia systems, network addressing, network interoperation, network optimization, network protocols, service creation

26 Migration: Optimizing the migration of virtual computers



Constantine P. Sapuntzakis, Ramesh Chandra, Ben Pfaff, Jim Chow, Monica S. Lam, Mendel Rosenblum

December 2002 **ACM SIGOPS Operating Systems Review**, Volume 36 Issue SI

Publisher: ACM Press

Full text available: pdf(1.68 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

This paper shows how to quickly move the state of a running computer across a network, including the state in its disks, memory, CPU registers, and I/O devices. We call this state a *capsule*. Capsule state is hardware state, so it includes the entire operating system as well as applications and running processes. We have chosen to move x86 computer states because x86 computers are common, cheap, run the software we use, and have tools for migration. Unfortunately, x86 c ...

27 An XML query engine for network-bound data

Zachary G. Ives, A. Y. Halevy, D. S. Weld

December 2002 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 11 Issue 4

Publisher: Springer-Verlag New York, Inc.

Full text available:  pdf(351.86 KB) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

XML has become the lingua franca for data exchange and integration across administrative and enterprise boundaries. Nearly all data providers are adding XML import or export capabilities, and standard XML Schemas and DTDs are being promoted for all types of data sharing. The ubiquity of XML has removed one of the major obstacles to integrating data from widely disparate sources - namely, the heterogeneity of data formats. However, general-purpose integration of data across the wide area is also re ...

Keywords: Data integration, Data streams, Query processing, Web and databases, XML

28 [Articles: The Deliberate Revolution](#)




Mike Burner

March 2003 **Queue**, Volume 1 Issue 1

Publisher: ACM Press

Full text available:  pdf(326.81 KB)

 html(64.11 KB)

Additional Information: [full citation](#), [index terms](#)

29 [Run-time adaptation in river](#)



Remzi H. Arpaci-Dusseau

February 2003 **ACM Transactions on Computer Systems (TOCS)**, Volume 21 Issue 1

Publisher: ACM Press

Full text available:  pdf(849.04 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We present the design, implementation, and evaluation of run-time adaptation within the River dataflow programming environment. The goal of the River system is to provide adaptive mechanisms that allow database query-processing applications to cope with performance variations that are common in cluster platforms. We describe the system and its basic mechanisms, and carefully evaluate those mechanisms and their effectiveness. In our analysis, we answer four previously unanswered and important que ...

Keywords: Performance availability, clusters, parallel I/O, performance faults, robust performance, run-time adaptation

30 [Industry track papers and presentations: technology trends: Building enterprise portals: principles to practice](#)



Tushar K. Hazra

May 2002 **Proceedings of the 24th International Conference on Software Engineering**

Publisher: ACM Press

Full text available:  pdf(1.85 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Primary objective of this paper is to offer an exclusive view of constructing and deploying enterprise portals by using a component-based development approach. As the dot-com hype dies down, most companies are forced to revisit their enterprise-wide Web integration strategies. This paper offers a pragmatic roadmap that these companies may follow in their upcoming enterprise portal deployment initiatives. The academic world plays a significant role in the advances of the portal technology. In this ...

31 [Business Continuity Planning](#)

Martin Nemzow

July 1997 **International Journal of Network Management**, Volume 7 Issue 3

Publisher: John Wiley & Sons, Inc.

Full text available:  pdf(295.36 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

This article considers various strategies for protecting an organisation from both natural and man-made disasters. The differences between business continuity planning, and disaster recovery planning are recognised. © 1997 John Wiley & Sons, Ltd.

32 Measurement tools: Robust synchronization of software clocks across the internet



Darryl Veitch, Satish Babu, Attila Pásztor

October 2004 **Proceedings of the 4th ACM SIGCOMM conference on Internet measurement**

Publisher: ACM Press

Full text available: pdf(1.59 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Accurate, reliable timestamping which is also convenient and inexpensive is needed in many important areas including real-time network applications and network measurement. Recently the TSC register, which counts CPU cycles in popular PC architectures, was proposed as the basis of a new software clock which in terms of rate performance performs as well as more expensive GPS alternatives. Smooth and precise clock rate is essential to measure time differences accurately. We show ...

Keywords: GPS, NTP, network measurement, round-trip time, software clock, synchronization, timing

33 Implementing a network improvement action plan



Kim Todd, Jon Rickman, Tabatha Verbick

October 2004 **Proceedings of the 32nd annual ACM SIGUCCS conference on User services**

Publisher: ACM Press

Full text available: pdf(182.43 KB) Additional Information: [full citation](#), [abstract](#), [index terms](#)

Firewall failures and the recent onslaught of computer viruses and worms, such as Klez, loveletter, SoBig, and BadBoy, increased network traffic and ever-growing network abuse, propelled security awareness at Northwest Missouri State University to a higher level and increased the need for finding and implementing effective solutions.

Northwest's Information Systems Department was able to build awareness of computer security issues, decrease network abuse, institute significant enhance ...

Keywords: communication, customer service, network management, network stability

Results 21 - 33 of 33

Result page: [previous](#) [1](#) [2](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads: [Adobe Acrobat](#) [QuickTime](#) [Windows Media Player](#) [Real Player](#)



USPTO

[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: ☒ The ACM Digital Library ☐ The Guide

"reclaim lock"

SEARCH

THE ASSOCIATION FOR COMPUTING MACHINERY


[Feedback](#) [Report a problem](#) [Satisfaction survey](#)
Terms used **reclaim lock**

Found 2 of 186,958

Sort results by

relevance

Display results

expanded form

Save results to a Binder

Search Tips

☐ Open results in a new window
Try an [Advanced Search](#)Try this search in [The ACM Guide](#)

Results 1 - 2 of 2

Relevance scale ☐ ☐ ☐ ☐ ☐

1 [Effective fine-grain synchronization for automatically parallelized programs using optimistic synchronization primitives](#)



Martin C. Rinard

November 1999 **ACM Transactions on Computer Systems (TOCS)**, Volume 17 Issue 4

Publisher: ACM Press

Full text available: pdf(637.69 KB)

 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

This article presents our experience using optimistic synchronization to implement fine-grain atomic operations in the context of a parallelizing compiler for irregular, object-based computations. Our experience shows that the synchronization requirements of these programs differ significantly from those of traditional parallel computations, which use loop nests to access dense matrices using affine access functions. In addition to coarse-grain barrier synchronization, our irregular comput ...

Keywords: atomic operations commutativity analysis, optimistic synchronization, parallel computing, parallelizing compilers, synchronization

2 [Architectural principles and techniques for distributed multimedia application support in operating systems](#)



Geoff Coulson, Gordon Blair

October 1995 **ACM SIGOPS Operating Systems Review**, Volume 29 Issue 4

Publisher: ACM Press

Full text available: pdf(659.83 KB)

 Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

We propose some architectural principles we have found useful for the support of continuous media applications in a microkernel environment. In particular, we discuss i) the principle of upcall-driven application structuring whereby communications events are system rather than application initiated, ii) the principle of split-level system structuring whereby key system functions are carried out co-operatively between kernel and user level components and iii) the principle of decoupling of contro ...

Results 1 - 2 of 2

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

 Useful downloads: [Adobe Acrobat](#) [QuickTime](#) [Windows Media Player](#) [Real Player](#)


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: ☒ The ACM Digital Library ☐ The Guide

frangipani +lock +reclaim

SEARCH

THE ACM DIGITAL LIBRARY


[Feedback](#) [Report a problem](#) [Satisfaction survey](#)
Terms used **frangipani lock reclaim**

Found 6 of 37 searched out of 37.

Sort results by

relevance ☒

Display results

expanded form ☒

Save results to a Binder

Search Tips

☐ Open results in a new windowTry an [Advanced Search](#)Try this search in [The ACM Guide](#)

Results 1 - 6 of 6

Relevance scale ☐ ☐ ☐ ☐ ☐1 [Frangipani: a scalable distributed file system](#)

Chandramohan A. Thekkath, Timothy Mann, Edward K. Lee

October 1997 **ACM SIGOPS Operating Systems Review**, **Proceedings of the sixteenth ACM symposium on Operating systems principles SOSP '97**, Volume 31 Issue 5

Publisher: ACM Press

Full text available: pdf(2.20 MB)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)2 [Decentralized storage systems: Ivy: a read/write peer-to-peer file system](#)

Athicha Muthitacharoen, Robert Morris, Thomer M. Gil, Benjie Chen

December 2002 **ACM SIGOPS Operating Systems Review**, Volume 36 Issue SI

Publisher: ACM Press

Full text available: pdf(1.65 MB)

Additional Information: [full citation](#), [abstract](#), [references](#)

Ivy is a multi-user read/write peer-to-peer file system. Ivy has no centralized or dedicated components, and it provides useful integrity properties without requiring users to fully trust either the underlying peer-to-peer storage system or the other users of the file system. An Ivy file system consists solely of a set of logs, one log per participant. Ivy stores its logs in the DHash distributed hash table. Each participant finds data by consulting all logs, but performs modifications by appendi ...

3 [File and storage systems: The Google file system](#)

Sanjay Ghemawat, Howard Gobioff, Shun-Tak Leung

October 2003 **Proceedings of the nineteenth ACM symposium on Operating systems principles**

Publisher: ACM Press

Full text available: pdf(275.54 KB)

Additional Information: [full citation](#), [references](#), [index terms](#)**Keywords:** clustered storage, data storage, fault tolerance, scalability4 [Decentralized storage systems: Farsite: federated, available, and reliable storage for an incompletely trusted environment](#)


Atul Adya, William J. Bolosky, Miguel Castro, Gerald Cermak, Ronnie Chaiken, John R. Douceur, Jon Howell, Jacob R. Lorch, Marvin Theimer, Roger P. Wattenhofer

December 2002 **ACM SIGOPS Operating Systems Review**, Volume 36 Issue SI

Publisher: ACM Press

Full text available:

Additional Information:

 pdf(1.87 MB)[full citation](#), [abstract](#), [references](#)

Farsite is a secure, scalable file system that logically functions as a centralized file server but is physically distributed among a set of untrusted computers. Farsite provides file availability and reliability through randomized replicated storage; it ensures the secrecy of file contents with cryptographic techniques; it maintains the integrity of file and directory data with a Byzantine-fault-tolerant protocol; it is designed to be scalable by using a distributed hint mechanism and delegatio ...

5 [The evolution of Coda](#)



M. Satyanarayanan

May 2002 **ACM Transactions on Computer Systems (TOCS)**, Volume 20 Issue 2

Publisher: ACM Press

Full text available:  pdf(441.35 KB)Additional Information: [full citation](#), [abstract](#), [references](#), [citings](#), [index terms](#)

Failure-resilient, scalable, and secure read-write access to shared information by mobile and static users over wireless and wired networks is a fundamental computing challenge. In this article, we describe how the Coda file system has evolved to meet this challenge through the development of mechanisms for server replication, disconnected operation, adaptive use of weak connectivity, isolation-only transactions, translucent caching, and opportunistic exploitation of hardware surrogates. For eac ...

Keywords: Adaptation, Linux, UNIX, Windows, caching, conflict resolution, continuous data access, data staging, disaster recovery, disconnected operation, failure, high availability, hoarding, intermittent networks, isolation-only transactions, low-bandwidth networks, mobile computing, optimistic replica control, server replication, translucent cache management, weakly connected operation


6 [Decentralized storage systems: Taming aggressive replication in the Pangaea wide-area file system](#)



Yasushi Saito, Christos Karamanolis, Magnus Karlsson, Mallik Mahalingam

December 2002 **ACM SIGOPS Operating Systems Review**, Volume 36 Issue SI

Publisher: ACM Press

Full text available:  pdf(1.93 MB)Additional Information: [full citation](#), [abstract](#), [references](#)

Pangaea is a wide-area file system that supports data sharing among a community of widely distributed users. It is built on a symmetrically decentralized infrastructure that consists of commodity computers provided by the end users. Computers act autonomously to serve data to their local users. When possible, they exchange data with nearby peers to improve the system's overall performance, availability, and network economy. This approach is realized by aggressively creating a replica of a file w ...

Results 1 - 6 of 6

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)